International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 10, October 2026)

AI Career Navigator: An Explainable IT Career Readiness Platform with Hybrid Semantic Skill Alignment and Gap-to-Action Planning

Sumit Chaurasia[1], Prajwal Jagatap[2], Suyash Karandikar[3], Harsh Uttamani[4], Vrushali Shinde[5]
Department of MCA, P.E.S.'s Modern College of Engineering, Pune, India

***Abstract*— Career-guidance tools can combine occupational taxonomies, resume parsing, vacancy matching, proficiency assessment, and learning recommendations, yet their outputs may be difficult to audit. This paper presents AI Career Navigator, an implemented student-focused decision-support prototype that separates canonical skill normalization, exact and curated-alias matching, Sentence-BERT (SBERT) semantic comparison, proficiency-adjusted readiness, resume keyword overlap, and gap-to-action guidance. Its seeded catalogue contains 10 IT tracks, 50 weighted career-skill links, 30 canonical skills, and 32 curated aliases. The semantic matcher uses sentence-transformers/all-MiniLM-L6-v2 with cosine similarity and an initial threshold of 0.62. On a balanced, author-labelled pilot set of 64 requirement/evidence pairs, hybrid matching achieved 65.62% accuracy and an F1 score of 0.476, compared with 60.94% and 0.359 for exact/alias matching; recall remained 31.25%. A threshold sweep on the same examples reached an F1 score of 0.842 at 0.45, but this is exploratory and not an unbiased held-out estimate. Local functional verification also passed 538 backend tests with one skipped test and completed a production frontend build. Because labels are not independently adjudicated and the sample is not representative, the evidence supports prototype feasibility rather than generalizable recruitment accuracy or improved hiring outcomes.**

***Keywords*— Career readiness, explainable recommendation, hybrid matching, sentence embeddings, skill-gap analysis, student career guidance.**

I. Introduction

Mismatch between candidate capabilities and job requirements is not only a retrieval problem. Vacancies and resumes describe related skills using different terms, abbreviations, tools, and contexts. Exact keyword matching is easy to explain but may miss paraphrases, while embedding-only matching may treat related but non-interchangeable technologies as equivalent. Career-guidance applications also need to distinguish a vacancy-specific overlap score from a broader readiness estimate, then turn identified gaps into practical next steps.

Sentence-BERT (SBERT) provides sentence or phrase embeddings that can be compared using cosine similarity [1]. Occupational resources such as O*NET and ESCO provide structured occupation and skill terminology [2], [3], but their concepts and relationships do not map automatically into a small application-specific catalogue. Such mappings require curation and review. A system should therefore expose which matching route was used and make uncertainty and data limitations visible rather than presenting a single score as definitive evidence of competence.

AI Career Navigator is an implemented web application that links resume evidence, candidate profiles, selected job requirements, skill-gap analysis, and learning actions. Its applied contribution is the separation of evidence and score layers in an inspectable workflow: users can review matched and missing requirements, distinguish exact/alias results from semantic matches and keyword fallback, inspect the assumptions used for readiness, and interpret market indicators as dated context rather than a live vacancy or salary feed. This paper documents the implementation and an exploratory pilot comparison; it does not claim a new transformer architecture, a validated hiring predictor, or a placement guarantee.

II. Related Work

SBERT was introduced to produce sentence representations that can be compared efficiently, including through cosine similarity [1]. The pretrained all-MiniLM-L6-v2 model provides a compact sentence-embedding option for semantic comparison [10]. Its similarity scores are linguistic signals; they do not by themselves establish that two technologies are interchangeable, that a person has mastered a skill, or that resume evidence is reliable.

Occupational vocabularies provide a structured foundation for career-facing systems. O*NET describes occupations, skills, knowledge, and work activities [2], while ESCO links occupations with skills, competences, and qualifications [3].

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 10, October 2026)

These resources improve terminology and provenance, but registering a source does not establish that every concept has been mapped into an application catalogue or independently validated.

Career-guidance research has combined skill profiling, job-role exploration, and network visualization, as in C3-IoC [4], and has explored personalized guidance using resume parsing and user characteristics [5]. Resume-to-position research has applied BERT-based semantic representations to compare candidate and vacancy text [6]. Other systems connect career-roadmap milestones and progress tracking with recommendations [7], or combine skills mapping with mentoring and labor-market context [8]. A closely related hybrid approach combines semantic similarity and ontology-based matching for skill-gap and career-path recommendations [9].

These works show that semantic matching, skill-gap analysis, and job-informed learning recommendations are established directions. In particular, prior work has already combined semantic and ontology-based matching for career-path recommendations [9], so this paper does not claim that such a combination is new. Its narrower contribution is an applied, source-aware workflow that separates vacancy overlap, proficiency-adjusted readiness, and resume keyword coverage, while exposing the matching path and fallback state. A controlled evaluation using independently labelled data is still needed to determine whether its matching decisions improve over simpler baselines.

III. System Design And Methodology

The workflow begins with selected occupational metadata, curated IT tracks, and job-posting evidence. Skill labels are normalized into a canonical layer, candidate and vacancy phrases are compared through a hybrid matcher, and results are presented as distinct scores that can support prioritized learning actions. Figure 1 summarizes the data and decision flow.



Fig. 1. Source-Aware Candidate/Job Alignment and Readiness Workflow.

A. Canonical Skill Catalogue and Provenance

The seeded catalogue contains 10 IT tracks and 50 weighted career-skill relationships. The current prototype reports 30 canonical skills and 32 curated aliases. Four provenance labels are registered: O*NET, ESCO, NSDC, and prototype. Source-processing scripts produce review artefacts for unmatched or ambiguous terms. The source candidates and generated mappings must not be interpreted as a complete, adjudicated O*NET/ESCO ontology; the catalogue remains a curated application layer.

B. Resume and Candidate Evidence

Resume ingestion accepts searchable PDF, DOCX, and TXT files, extracts text, detects known skill labels and aliases, and produces canonical matches. Extracted terms can be associated with a signed-in candidate profile. Parser-derived and self-reported entries remain evidence records rather than independently verified competencies. PDF extraction depends on a readable text layer unless separate OCR support is configured, and an incomplete vocabulary may omit tools that appear in free text.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 10, October 2026)

C. Hybrid Exact/Alias and Semantic Matching

For each selected vacancy, the application detects technical requirements from provider tags, titles, and description patterns. In hybrid mode, exact or curated-alias matching is applied first. Remaining requirements are compared with unused candidate skill phrases after including short job-title or description context. The semantic component uses sentence-transformers/all-MiniLM-L6-v2 to encode the phrases and calculates cosine similarity [10]. Its initial decision threshold is 0.62, configured through SBERT_MATCH_THRESHOLD. A greedy assignment prevents one candidate phrase from satisfying multiple semantic requirements in the same match.

The threshold is a starting configuration, not a validated universal cutoff. If the model cannot load, the API identifies the result as keyword_fallback rather than presenting it as a semantic result. The matching-status endpoint reports dependency status, model-loaded state, active threshold, and loading errors. The browser result also indicates which path actually produced the match, allowing users to distinguish an embedding-based decision from a fallback result.

D. Score Definitions and Gap-to-Action Guidance

The vacancy match score is equal-weight skill coverage over the detected requirements: $J = 100 \times M/N$, where M is the number of requirements marked matched and N is the number of detected technical requirements. When $N = 0$, the score is unavailable rather than reported as 0% or 100%. It is not a hiring prediction.

Profile readiness is proficiency-adjusted coverage across the selected career requirements: $R = 100 \times [\Sigma_i \min(p_i, r_i)] / [\Sigma_i r_i]$, where p_i is normalized profile proficiency and r_i is the required proficiency on a common project scale. If no requirements are mapped, readiness is not scorable. Resume ATS overlap is calculated separately using keyword coverage. The interface keeps J, R, and ATS overlap separate because they describe different properties.

Detected missing or weak skills feed the action center, which presents prioritized learning topics, practical work, and study timelines where resource mappings exist. When there is insufficient learning history, the Readiness Velocity widget displays “Awaiting activity” rather than fabricating a weekly rate. The application also includes a dated India hiring-indicator snapshot and qualitative skill-focus context. These indicators are contextual, publisher-specific signals; they do not represent a complete vacancy count, role-level salary estimates, or an individual probability of being hired.

IV. Implementation And Data Provenance

The web application uses a React/Vite frontend, a Flask API, SQLAlchemy persistence, and PostgreSQL. The semantic dependency is the sentence-transformers library with the pretrained all-MiniLM-L6-v2 model [10]. The model is loaded lazily in the Flask process; a successful standalone warm-up does not by itself prove that a running web worker has loaded it. The API status endpoint and the result metadata are therefore used to confirm which matcher was active for an individual request.

Occupational provenance labels identify source families used in the project, but only the curated application catalogue is treated as active matching data. Unmatched and ambiguous source terms are retained for review instead of being silently forced into equivalence. This design supports traceability while acknowledging that crosswalking large taxonomies into a smaller career catalogue is a data-curation task.

Local smoke checks demonstrated model loading, the matching-status endpoint, the hybrid-mode interface, resume extraction, a five-requirement Software Developer example with two matched keywords (40% overlap), and generation of a four-page PDF report. These checks demonstrate that selected application paths execute; they are not estimates of model effectiveness, user benefit, or production security.

V. Evaluation Setup

The pilot benchmark contains 64 manually authored requirement/evidence pairs: 32 positive matches and 32 hard negatives. The negative examples include related but non-interchangeable technologies, such as Java versus JavaScript, AWS versus Azure, and Docker versus Kubernetes. The labels were created by the project authors and have not been independently adjudicated. The small balanced set is useful for checking matching behavior but does not represent the distribution of skills, resumes, or vacancies in a live recruitment setting.

Three methods were compared: (A) exact/alias matching, (B) SBERT cosine similarity at the configured 0.62 threshold, and (C) hybrid matching that applies exact/alias matching first and SBERT to unresolved pairs. The evaluation records accuracy, precision, recall, F1, specificity, confusion counts, pair-level similarity, and threshold sensitivity. The 0.62 result is the configured-threshold comparison. Threshold sweeps are diagnostic only because choosing a threshold and evaluating it on the same examples produces an optimistic estimate.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 10, October 2026)

The application also contains offline regression tests for matching control flow, fallback reporting, empty requirements, and reuse prevention. Stronger validation would require a larger set of consented and de-identified resume/job evidence, independent expert labels, inter-annotator agreement, a held-out set defined before threshold tuning, and uncertainty estimates. These requirements are not satisfied by the current author-labelled pilot.

VI. Results And Discussion

TABLE I. PILOT MATCHING METRICS AT THE CONFIGURED THRESHOLD (0.62)

Method	Accuracy	Precision	Recall	F1
Exact/alias	60.94%	100.00%	21.88%	0.359
SBERT	62.50%	100.00%	25.00%	0.400
Hybrid	**65.62%**	**100.00%**	**31.25%**	**0.476**

At the configured threshold, exact/alias matching achieved 60.94% accuracy, SBERT achieved 62.50%, and hybrid matching achieved 65.62%. Hybrid matching therefore improved accuracy by 4.68 percentage points and F1 by 11.72 percentage points over the exact/alias baseline on this set. The corresponding true-positive/false-positive/true-negative/false-negative counts were 7/0/32/25 for exact/alias, 8/0/32/24 for SBERT, and 10/0/32/22 for hybrid matching.

Hybrid recall was 31.25%, so 22 of the 32 positive pairs were missed at the configured threshold. Zero false positives among only 32 negative examples does not demonstrate perfect precision in a larger or differently distributed sample. An exact McNemar comparison found three pairs corrected by hybrid matching and none corrected only by the baseline; the two-sided p-value was 0.25, so the paired difference was not statistically significant.

TABLE II. HYBRID MATCHING THRESHOLD SENSITIVITY ON THE SAME 64 PAIRS

Threshold	Accuracy	Precision	Recall	F1
0.62	**65.62%**	**100.00%**	**31.25%**	**0.476**
0.50	81.25%	95.45%	65.62%	0.778
0.45	85.94%	96.00%	75.00%	0.842

The threshold sweep shows that the pilot results are highly threshold-sensitive. On the same 64 pairs, hybrid matching at 0.50 produced 81.25% accuracy, 95.45% precision, 65.62% recall, and 0.778 F1; at 0.45 it produced 85.94% accuracy, 96.00% precision, 75.00% recall, and 0.842 F1. These values are exploratory sensitivity results, not unbiased estimates of tuned performance, because threshold selection and evaluation used the same examples. A separate held-out set is required before selecting an operating threshold.

The full local backend regression suite reported 538 passed tests, one skipped test, and no failures; a focused regression run reported 94 passed tests, one skipped, and no failures. The frontend production build completed after transforming 1,599 modules. These results support local software functionality and regression coverage. They do not establish that the matching scores are valid measures of competence or that recommendations improve learning or employment outcomes.

VII. Explainability, Responsible Use, And Limitations

The interface reports the matching route actually used and does not label keyword fallback as semantic matching. It separates vacancy skill overlap, readiness based on profile proficiency and role requirements, and resume keyword coverage. These measures should be treated as decision-support signals, not proof of competence or probabilities of employment. Parser output may be incomplete, and self-reported skills should remain editable and should not be described as verified without supporting evidence.

The active career catalogue is small compared with O*NET and ESCO, and source registration does not mean that all source concepts have been mapped or validated. The configured threshold of 0.62 produced low recall on the pilot, while lowering the threshold increased recall but also allowed false positives. Semantic similarity captures related language rather than technology equivalence, proficiency, or evidence quality. The balanced, manually authored set, absent independent adjudication, and lack of a controlled user study prevent population-level generalization or claims of user benefit.

The market indicators are dated and publisher-specific. They cannot establish complete current vacancy counts, role-specific salaries, or individual hiring chances; salary values should remain unavailable until a verified India-specific salary source and as-of date are configured.

Before deployment, the project also requires security hardening: local testing flagged a 23-byte JWT HMAC key below the recommended 32-byte minimum for SHA-256, and HTTPS, access-control testing, email-verification settings, and software warnings require review. Candidate resumes and evaluation examples should be de-identified where appropriate, collected with consent, and access-restricted.

Further work should independently annotate a larger and more representative evidence set, measure inter-annotator agreement, define a held-out test set before threshold tuning, report uncertainty estimates, and review errors by skill family. A separate user evaluation is needed to test whether the learning actions are useful. These steps should precede claims about general performance, effectiveness, or readiness for production use.

VIII. Conclusion

AI Career Navigator implements a traceable workflow connecting canonical skill normalization, resume extraction, hybrid job-skill matching, proficiency-adjusted readiness, prioritized learning actions, and contextual hiring signals. It keeps vacancy overlap, profile readiness, and resume keyword coverage distinct and makes the matching path and fallback state inspectable. On an author-labelled 64-pair pilot, hybrid matching improved accuracy from 60.94% to 65.62% and F1 from 0.359 to 0.476 relative to exact/alias matching at the configured 0.62 threshold, but recall remained 31.25% and the paired difference was not statistically significant. The threshold sweep indicates sensitivity rather than a validated operating point. The current evidence supports the feasibility of an explainable prototype, not generalizable recruitment performance. Independent labels, a held-out evaluation, uncertainty reporting, and deployment security review remain necessary.

References

[1] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," in Proc. 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, 2019, pp. 3982–3992, doi: 10.18653/v1/D19-1410.

[2] National Center for O*NET Development, "O*NET Database Releases Archive: Release 31.0 (August 2026)," O*NET Resource Center, 2026. [Online]. Available: https://www.onetcenter.org/db_releases.html. Accessed: Oct. 18, 2026.

[3] European Commission, "About ESCO: European Skills, Competences, Qualifications and Occupations," ESCO v1.2.1, updated Dec. 10, 2025. [Online]. Available: https://esco.ec.europa.eu/en/about-esco. Accessed: Oct. 18, 2026.

[4] A. José-García, A. Sneyd, A. Melro, A. Ollagnier, G. Tarling, H. Zhang, M. Stevenson, R. Everson, and R. Arthur, "C3-IoC: A Career Guidance System for Assessing Student Skills using Machine Learning and Network Visualisation," International Journal of Artificial Intelligence in Education, vol. 33, pp. 1092–1119, 2023, doi: 10.1007/s40593-022-00317-y.

[5] D. P. Agarkar, R. Mahant, S. Chavan, and A. Bhosale, "A Comprehensive Base Paper: AI-Powered Personalized Career Guidance System," International Journal of Ingenious Research, Invention and Development, vol. 5, no. 3, June 2026, doi: 10.5281/zenodo.20509921.

[6] X. Liu, "Deep Learning-Based Intelligent Resume-Position Matching System: Semantic Understanding and Recommendation of BERT Model in Massive Recruitment Data," in Proc. 2025 International Symposium on Machine Learning and Social Computing (MLSC 2025), Hong Kong, China, Oct. 16–18, 2025, pp. 8–13, doi: 10.1145/3778450.3778452.

[7] R. Sreevatsan, J. Syed Raffi Ahamed, and A. B. Hajira Be, "Interactive Career Roadmap Website: Career Paths with Progress Tracking and AI-Guided Recommendations," International Journal of Creative and Open Research in Engineering and Management, vol. 2, no. 6, June 2026, doi: 10.55041/ijcope.v2i6.169.

[8] R. Wahrini, H. Hasbi, M. Nuruzzaman, N. A. Nur, P. Budiastuti, M. Gaspersz, and A. T. Mara, "AI-Driven Career Guidance to Reduce Vocational Students' Career Path Anxiety through Skills Mapping, Adaptive Mentoring, and Labor Market Intelligence," F1000Research, vol. 15, p. 298, 2026, doi: 10.12688/f1000research.174858.1.

[9] V. Abishek Narayan and A. Rajalakshmi, "AI-Powered Skill Gap Analysis and Career Path Recommendation System Using Hybrid Semantic and Ontology-Based Matching," in Proc. 2026 International Conference on Data Science and Business Systems (ICDSBS), Singapore, May 2026, doi: 10.1109/ICDSBS69077.2026.11634695.

[10] Sentence Transformers, "sentence-transformers/all-MiniLM-L6-v2: Model Card," Hugging Face. [Online]. Available: https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2. Accessed: Oct. 22, 2026.